

Lecture 7
Logical Agents
Inference in First Order Logic

Marco Chiarandini

Department of Mathematics & Computer Science
University of Southern Denmark

Slides by Stuart Russell and Peter Norvig

Summary

First-order logic:

- objects and relations are semantic primitives
- syntax: constants, functions, predicates, equality, quantifiers

Increased expressive power: sufficient to define wumpus world

Situation calculus:

- conventions for describing actions and change in FOL
- can formulate planning as inference on a situation calculus KB

Course Overview

- ✓ Introduction
 - ✓ Artificial Intelligence
 - ✓ Intelligent Agents
- ✓ Search
 - ✓ Uninformed Search
 - ✓ Heuristic Search
- ✓ Adversarial Search
 - ✓ Minimax search
 - ✓ Alpha-beta pruning
- Knowledge representation and Reasoning
 - ✓ Propositional logic
 - ✓ First order logic
 - Inference
- Uncertain knowledge and Reasoning
 - Probability and Bayesian approach
 - Bayesian Networks
 - Hidden Markov Chains
 - Kalman Filters
- Learning
 - Decision Trees
 - Maximum Likelihood
 - EM Algorithm
 - Learning Bayesian Networks
 - Neural Networks
 - Support vector machines

2

Outline

- ◇ Reducing first-order inference to propositional inference
- ◇ Unification
- ◇ Generalized Modus Ponens
- ◇ Forward and backward chaining
- ◇ Logic programming
- ◇ Resolution

A brief history of reasoning

450b.c.	Stoics	propositional logic, inference (maybe)
322b.c.	Aristotle	“syllogisms” (inference rules), quantifiers
1565	Cardano	probability theory (propositional logic + uncertainty)
1847	Boole	propositional logic (again)
1879	Frege	first-order logic
1922	Wittgenstein	proof by truth tables
1930	Gödel	\exists complete algorithm for FOL
1930	Herbrand	complete algorithm for FOL (reduce to propositional)
1931	Gödel	$\neg\exists$ complete algorithm for arithmetic
1960	Davis/Putnam	“practical” algorithm for propositional logic
1965	Robinson	“practical” algorithm for FOL—resolution

5

Definitions

For a predicate calculus expression X and an interpretation I :

- If X has a value of T under I and a particular variable assignment, then I is said to **satisfy** X .
- If I satisfies X for all variable assignments, then I is a **model** of X
- X is **satisfiable** if and only if there exist an interpretation and variable assignment that satisfy it; otherwise, it is **unsatisfiable**
- If a set of expressions is not satisfiable, it is said to be **inconsistent**
- If X has a value T for all possible interpretations, X is said to be **valid**.
Eg.: $(p(X) \wedge \neg p(X))$ while $\exists X(P(X) \vee \neg p(X))$

6

Rules of Inference for Propositions

Definition

A **Proof Procedure** is a combination of an **inference rule** and an algorithm for applying that rule to a set of logical expressions to generate new sentences.

Eg: **Resolution inference** rule.

Definition

A predicate calculus expression X **logically follows** from a set S of predicate calculus expressions if every interpretation and variable assignment that satisfies S also satisfies X .

An inference rule is **sound** if every predicate calculus expression produced by the rule from a set S of predicate calculus expressions also logically follows from S .

An inference rule is **complete** if, given a set S of predicate calculus expressions, the rule can infer every expression that logically follows from S .

7

Rule of inference	Name	Rule of inference	Name
$\frac{p}{p \rightarrow q}$ $\therefore q$	Modus Ponens	$\frac{p}{p \vee q}$	Addition
$\frac{\neg q}{p \rightarrow q}$ $\therefore \neg p$	Modus tollens	$\frac{p \vee q}{p}$	Simplification
$\frac{p \rightarrow q}{q \rightarrow r}$ $\therefore p \rightarrow r$	Hypothetical syllogism	$\frac{p}{p \vee q}$	Conjunction
$\frac{p \vee q}{\neg p}$ $\therefore q$	Disjunctive syllogism	$\frac{p \vee q}{\neg p \vee r}$ $\therefore q \vee r$	Resolution

8

Rules of Inference for Quantified Statements

Rule of inference

Name

$$\therefore \frac{\forall x P(x)}{P(c)}$$

Universal instantiation

$$\therefore \frac{P(c) \text{ for an arbitrary } c}{\forall x P(x)}$$

Universal generalization

$$\therefore \frac{\exists x P(x)}{P(c) \text{ for some element } c}$$

Existential instantiation

$$\therefore \frac{P(c) \text{ for some element } c}{\exists x P(x)}$$

Existential generalization

Universal instantiation (UI)

Every instantiation of a universally quantified sentence α is entailed by it:

$$\therefore \frac{\forall v \alpha}{\text{Subst}(\{v/c\}, \alpha)}$$

for any variable v and ground term c . (Note, here we used prolog notation.)

E.g., $\forall x \text{ King}(x) \wedge \text{Greedy}(x) \implies \text{Evil}(x)$ yields

$$\begin{aligned} & \text{King}(\text{John}) \wedge \text{Greedy}(\text{John}) \implies \text{Evil}(\text{John}) \\ & \text{King}(\text{Richard}) \wedge \text{Greedy}(\text{Richard}) \implies \text{Evil}(\text{Richard}) \\ & \text{King}(\text{Father}(\text{John})) \wedge \text{Greedy}(\text{Father}(\text{John})) \implies \text{Evil}(\text{Father}(\text{John})) \\ & \vdots \end{aligned}$$

9

10

Existential instantiation (EI)

For any sentence α , variable v , and constant symbol k
that does not appear elsewhere in the knowledge base:

$$\therefore \frac{\exists v \alpha}{\text{Subst}(\{v/k\}, \alpha)}$$

E.g., $\exists x \text{ Crown}(x) \wedge \text{OnHead}(x, \text{John})$ yields

$$\text{Crown}(C_1) \wedge \text{OnHead}(C_1, \text{John})$$

provided C_1 is a new constant symbol, called a **Skolem constant**

Another example: from $\exists x d(x^y)/dy = x^y$ we obtain

$$d(e^y)/dy = e^y$$

provided e is a new constant symbol

Existential instantiation contd.

UI can be applied several times to **add** new sentences;
the new KB is logically equivalent to the old

EI can be applied once to **replace** the existential sentence;
the new KB is **not** equivalent to the old,
but is satisfiable iff the old KB was satisfiable

Reduction to propositional inference

Suppose the KB contains just the following:

$$\forall x \text{ King}(x) \wedge \text{Greedy}(x) \implies \text{Evil}(x)$$

King(John)
Greedy(John)
Brother(Richard, John)

Instantiating the universal sentence in **all possible** ways, we have

$$\text{King}(\text{John}) \wedge \text{Greedy}(\text{John}) \implies \text{Evil}(\text{John})$$

$$\text{King}(\text{Richard}) \wedge \text{Greedy}(\text{Richard}) \implies \text{Evil}(\text{Richard})$$

King(John)
Greedy(John)
Brother(Richard, John)

The new KB is **propositionalized**: proposition symbols are

King(John), *Greedy(John)*, *Evil(John)*, *King(Richard)*, etc.

and can therefore be solved by the methods seen with **propositional logic**

13

Problems with propositionalization

Propositionalization seems to generate lots of irrelevant sentences.

E.g., from

$$\forall x \text{ King}(x) \wedge \text{Greedy}(x) \implies \text{Evil}(x)$$

King(John)
 $\forall y \text{ Greedy}(y)$
Brother(Richard, John)

it seems obvious that *Evil(John)*, but propositionalization produces lots of facts such as *Greedy(Richard)* that are irrelevant

With p k -ary predicates and n constants, there are $p \cdot n^k$ instantiations

With function symbols, it gets much much worse!

15

Reduction to propositional inference (contd.)

- Claim: a ground sentence is entailed by new KB iff entailed by original KB
- Claim: every FOL KB can be propositionalized so as to preserve entailment
- Idea: propositionalize KB and query, apply resolution, return result
- Problem: with variables and function symbols, there are infinitely many ground terms,
e.g., *Father(Father(Father(John)))*
- Theorem: Herbrand (1930). If a sentence α is entailed by an FOL KB, it is entailed by a **finite** subset of the propositional KB
- Idea: For $n = 0$ to ∞ do
create a propositional KB by instantiating with depth- n terms
see if α is entailed by this KB
- Problem: works if α is entailed, loops if α is not entailed
- Theorem: Turing (1936), Church (1936), entailment in FOL is **semidecidable**

14

Unification

We can get the inference immediately if we can find a substitution σ such that *King(x)* and *Greedy(x)* match *King(John)* and *Greedy(y)*

$\sigma = \{x/\text{John}, y/\text{John}\}$ works

$\text{Unify}(\alpha, \beta) = \sigma$ if $\alpha\sigma = \beta\sigma$

p	q	σ
<i>Knows(John, x)</i>	<i>Knows(John, Jane)</i>	$\{x/\text{Jane}\}$
<i>Knows(John, x)</i>	<i>Knows(y, OJ)</i>	$\{x/\text{OJ}, y/\text{John}\}$
<i>Knows(John, x)</i>	<i>Knows(y, Mother(y))</i>	$\{y/\text{John}, x/\text{Mother(John)}\}$
<i>Knows(John, x)</i>	<i>Knows(x, OJ)</i>	<i>fail</i>

Standardizing apart: rename variables to eliminate name overlap, e.g.,
Knows(z, OJ)

16

Generalized Modus Ponens (GMP)

Any inference in FOL has to use unification
 Here is an inference rule with the use of unification

$$\frac{p_1', p_2', \dots, p_n'}{(p_1 \wedge p_2 \wedge \dots \wedge p_n \Rightarrow q)} \quad \text{where } p_i'\sigma = p_i\sigma \text{ for all } i$$

$$\therefore q\sigma$$

p_1' is *King(John)* p_1 is *King(x)*
 p_2' is *Greedy(y)* p_2 is *Greedy(x)*
 σ is $\{x/\text{John}, y/\text{John}\}$ q is *Evil(x)*
 $q\sigma$ is *Evil(John)*

GMP used with KB of **definite clauses** (**exactly** one positive literal)
 All variables assumed universally quantified

17

Unification

Unification: search substitution that match two expressions

- constants (ground instances) cannot be substituted
- only variables can be substituted
- cannot substitute x by $p(x) \rightsquigarrow$ creates infinite regression
 occur check
- a variable can be substituted with another variable
- future substitutions must be consistent (substitution sequence)

Composition of substitutions:

$$\{Y/X, Z/W\}; \{X/V\}; \{V/a, W/f(b)\}$$

19

Soundness of GMP

Need to show that

$$p_1', \dots, p_n', (p_1 \wedge \dots \wedge p_n \Rightarrow q) \models q\sigma$$

provided that $p_i'\sigma = p_i\sigma$ for all i

Lemma: For any definite clause p , we have $p \models p\sigma$ by UI

1. $(p_1 \wedge \dots \wedge p_n \Rightarrow q) \models (p_1 \wedge \dots \wedge p_n \Rightarrow q)\sigma = (p_1\sigma \wedge \dots \wedge p_n\sigma \Rightarrow q\sigma)$
2. $p_1', \dots, p_n' \models p_1' \wedge \dots \wedge p_n' \models p_1'\sigma \wedge \dots \wedge p_n'\sigma$
3. From 1 and 2, $q\sigma$ follows by ordinary Modus Ponens

18

Unification

Unifiers must be as general as possible otherwise eliminate possibility for future solutions:

Eg: $p(X), p(Y)$ and $\{X/\text{fred}, Y/\text{fred}\}$

Definition

If μ is any unifier of expressions E and σ is a **most general unifier** then for μ applied to E there exists μ' such that $E\sigma = E\sigma\mu'$ where $E\mu$ and $E\sigma\mu'$ is the composition of unifiers.

mgu is unique (except for relabelling)

20

Example knowledge base contd.

... it is a crime for an Dane to sell weapons to hostile nations:

$Dane(x) \wedge Weapon(y) \wedge Sells(x, y, z) \wedge Hostile(z) \implies Criminal(x)$

Nono ... has some missiles, i.e., $\exists x Owns(Nono, x) \wedge Missile(x)$:

$Owns(Nono, M_1)$ and $Missile(M_1)$

... all of its missiles were sold to it by Colonel Thor

$\forall x Missile(x) \wedge Owns(Nono, x) \implies Sells(Thor, x, Nono)$

Missiles are weapons:

$Missile(x) \implies Weapon(x)$

An enemy of Denmark counts as "hostile":

$Enemy(x, Denmark) \implies Hostile(x)$

Thor, who is Dane ...

$Dane(Thor)$

The country Nono, an enemy of Denmark ...

$Enemy(Nono, Denmark)$

Forward chaining algorithm


```

function FOL-FC-Ask(KB,  $\alpha$ ) returns a substitution or false
  repeat until new is empty
    new  $\leftarrow$  {}
    for each sentence r in KB do
      ( $p_1 \wedge \dots \wedge p_n \implies q$ )  $\leftarrow$  Standardize-Apart(r)
      for each  $\delta$  such that ( $p_1 \wedge \dots \wedge p_n$ ) $\delta$  = ( $p'_1 \wedge \dots \wedge p'_n$ ) $\delta$ 
        for some  $p'_1, \dots, p'_n$  in KB
           $q' \leftarrow$  Subst( $\delta, q$ )
          if  $q'$  is not a renaming of a sentence already in KB or new
            then do
              add  $q'$  to new
               $\sigma \leftarrow$  Unify( $q', \alpha$ )
              if  $\sigma$  is not fail then return  $\sigma$ 
  add new to KB
  return false
  
```

25

26

Forward chaining proof



27

Properties of forward chaining

For first-order definite clauses the algorithm is:

- Sound because application of generalized modus ponens
- Complete (proof similar to propositional proof)

Datalog = first-order definite clauses + **no functions** (e.g., crime KB)

FC terminates for Datalog in poly iterations: at most $p \cdot n^k$ literals

May not terminate in general (with functions) if α is not entailed

This is unavoidable: entailment with definite clauses is semidecidable


28

Efficiency of forward chaining

- Simple observation: no need to match a rule on iteration k if a premise wasn't added on iteration $k - 1$
 \implies match each rule whose premise contains a newly added literal
- Matching itself can be expensive:
 - Database indexing allows $O(1)$ retrieval of known facts
 e.g., query $Missile(x)$ retrieves $Missile(M_1)$
 - Matching conjunctive premises against known facts is NP-hard

Forward chaining is widely used in [deductive databases](#)

Hard matching example



$$\begin{aligned}
 &Diff(wa, nt) \wedge Diff(wa, sa) \wedge \\
 &Diff(nt, q) Diff(nt, sa) \wedge \\
 &Diff(q, nsw) \wedge Diff(q, sa) \wedge \\
 &Diff(nsw, v) \wedge Diff(nsw, sa) \wedge \\
 &Diff(v, sa) \implies Colorable()
 \end{aligned}$$

$$\begin{aligned}
 &Diff(Red, Blue) \quad Diff(Red, Green) \\
 &Diff(Green, Red) \quad Diff(Green, Blue) \\
 &Diff(Blue, Red) \quad Diff(Blue, Green)
 \end{aligned}$$

$Colorable()$ is inferred iff the CSP has a solution
 CSPs include 3SAT as a special case, hence matching is NP-hard

29

30

Backward chaining algorithm

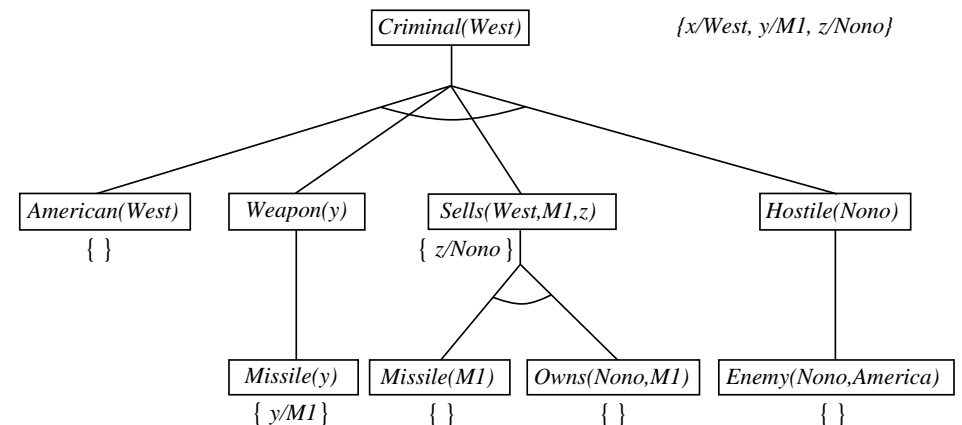
```

function FOL-BC-Ask(KB, goals,  $\sigma$ ) returns a set of substitutions
inputs: KB, a knowledge base
           goals, a list of conjuncts forming a query ( $\sigma$  already applied)
            $\sigma$ , the current substitution, initially the empty substitution {}
local variables: answers, a set of substitutions, initially empty

if goals is empty then return { $\sigma$ }
 $q' \leftarrow \text{Subst}(\sigma, \text{First}(\text{goals}))$ 
for each sentence r in KB
    where Standardize-Apart(r) = ( $p_1 \wedge \dots \wedge p_n \Rightarrow q$ )
    and  $\sigma' \leftarrow \text{Unify}(q, q')$  succeeds
     $\text{new\_goals} \leftarrow [p_1, \dots, p_n | \text{Rest}(\text{goals})]$ 
     $\text{answers} \leftarrow \text{FOL-BC-Ask}(\text{KB}, \text{new\_goals}, \text{Compose}(\sigma', \sigma)) \cup \text{answers}$ 
return answers
    
```

31

Backward chaining example



32

Properties of backward chaining

Depth-first recursive proof search: space is linear in size of proof

Incomplete due to infinite loops

⇒ fix by checking current goal against every goal on stack

Inefficient due to repeated subgoals (both success and failure)

⇒ fix using caching of previous results (extra space!)

Widely used (without improvements!) for **logic programming**

Logic programming

Sound bite: computation as inference on logical KBs

Logic programming

1. Identify problem
2. Assemble information
3. Tea break
4. Encode information in KB
5. Encode problem instance as facts
6. Ask queries
7. Find false facts

Ordinary programming

1. Identify problem
2. Assemble information
3. Figure out solution
4. Program solution
5. Encode problem instance as data
6. Apply program to data
7. Debug procedural errors

Should be easier to debug *Capital(NewYork,US)* than $x := x + 2$!

33

34

Prolog systems

Basis: backward chaining with Horn clauses + bells & whistles

Program = set of clauses = head :- literal₁, ... literal_n.

criminal(X) :- american(X), weapon(Y), sells(X,Y,Z), hostile(Z)

Efficient unification by **open coding**

Efficient retrieval of matching clauses by direct linking

Depth-first, left-to-right backward chaining

Built-in predicates for arithmetic etc., e.g., X is Y*Z+3

Closed-world assumption ("negation as failure")

e.g., given alive(X) :- not dead(X).

alive(joe) succeeds if dead(joe) fails

Resolution: brief summary

Full first-order version:

$$\frac{l_1 \vee \dots \vee l_k \quad m_1 \vee \dots \vee m_n}{\therefore (l_1 \vee \dots \vee l_{i-1} \vee l_{i+1} \vee \dots \vee l_k \vee m_1 \vee \dots \vee m_{j-1} \vee m_{j+1} \vee \dots \vee m_n)\sigma}$$

where $\text{Unify}(l_i, \neg m_j) = \sigma$.

For example,

$$\frac{\neg \text{Rich}(x) \vee \text{Unhappy}(x) \quad \text{Rich}(\text{Ken})}{\text{Unhappy}(\text{Ken})}$$

with $\sigma = \{x/\text{Ken}\}$

↔ Apply resolution steps to $\text{CNF}(KB \wedge \neg\alpha)$; complete for FOL

35

36

Conversion to CNF

Everyone who loves all animals is loved by someone:

$$\forall x [\forall y \text{ Animal}(y) \implies \text{Loves}(x, y)] \implies [\exists y \text{ Loves}(y, x)]$$

1. Eliminate biconditionals and implications

$$\forall x [\neg \forall y \neg \text{Animal}(y) \vee \text{Loves}(x, y)] \vee [\exists y \text{ Loves}(y, x)]$$

2. Move \neg inwards: $\neg \forall x, p \equiv \exists x \neg p$, $\neg \exists x, p \equiv \forall x \neg p$:

$$\forall x [\exists y \neg(\neg \text{Animal}(y) \vee \text{Loves}(x, y))] \vee [\exists y \text{ Loves}(y, x)]$$

$$\forall x [\exists y \neg \neg \text{Animal}(y) \wedge \neg \text{Loves}(x, y)] \vee [\exists y \text{ Loves}(y, x)]$$

$$\forall x [\exists y \text{ Animal}(y) \wedge \neg \text{Loves}(x, y)] \vee [\exists y \text{ Loves}(y, x)]$$

Conversion to CNF contd.

3. Standardize variables: each quantifier should use a different one

$$\forall x [\exists y \text{ Animal}(y) \wedge \neg \text{Loves}(x, y)] \vee [\exists z \text{ Loves}(z, x)]$$

4. Skolemize: a more general form of existential instantiation. Each existential variable is replaced by a **Skolem function** of the enclosing universally quantified variables:

$$\forall x [\text{Animal}(F(x)) \wedge \neg \text{Loves}(x, F(x))] \vee \text{Loves}(G(x), x)$$

5. Drop universal quantifiers:

$$[\text{Animal}(F(x)) \wedge \neg \text{Loves}(x, F(x))] \vee \text{Loves}(G(x), x)$$

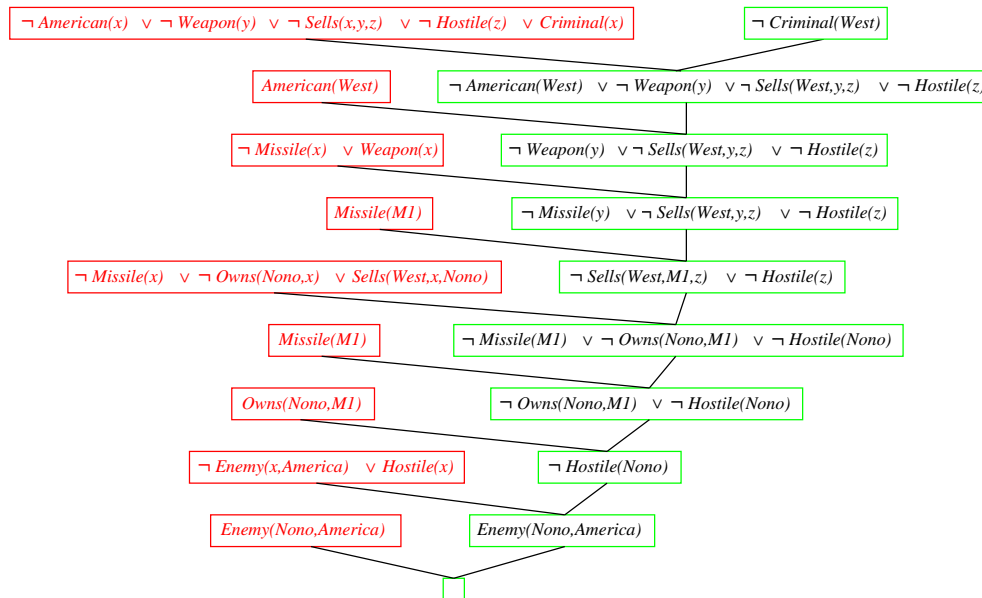
6. Distribute \wedge over \vee :

$$[\text{Animal}(F(x)) \vee \text{Loves}(G(x), x)] \wedge [\neg \text{Loves}(x, F(x)) \vee \text{Loves}(G(x), x)]$$

37

38

Resolution proof: definite clauses



39